"BBN Job #41540"
Bob Kahn Interview
August 10, 1994, Tape Three
-- CARBINER GROUP

QUESTION

BOB KAHN

Okay, I'm Bob Kahn, I'm currently the president of the Corporation for National Research Initiatives. Uh, in 1969, I was, uh, the senior scientist at Loprack(?) and Neumann(?), and I had, uh, just begun working on the Arpinet Project.

QUESTION

BOB KAHN

I'm going to go back to your original question.

QUESTION

BOB KAHN

My name is Bob Kahn. I'm currently President of the Corporation for National Research Initiatives,

was ...

QUESTION

BOB KAHN

You mean, the camera, it's no good?

QUESTION

BOB KAHN

Okay. Okay.

QUESTION

BOB KAHN

Okay, let's do it again.

QUESTION

BOB KAHN

All right.

QUESTION

BOB KAHN

Oh, I thought you wanted it on the record, on the tape.

QUESTION

BOB KAHN

Oh, so ...

QUESTION

BOB KAHN

Oh, so, we don't need to go back over that again.

QUESTION

BOB KAHN

All right, so, then, let's forget it.

QUESTION

BOB KAHN

I can't tell.

QUESTION

BOB KAHN

Uh, the story of the Open Net(?)?

QUESTION

BOB KAHN

Well, uh, I joined BB&N in 1966. And at the time I'd been a professor at MIT in the Electrical Engineering Department. And I'd taken a leave of absence with the idea that I would get involved in some practical kinds of activities.

And the one I happened to choose to work on was computer networking, probably with the encouragement of several of the people at BB&N. I'd started to develop some notions of what computer networking might be like, and sort of to document it. And, uh, by the time we had actually gotten the Arpinet(?) Award, I had decided to actually get involved in the implementation of that as well. I was principally involved in the system design aspects before that, but I eventually decided that, uh, being actually part of the project itself was probably the best way to get this practical experience.

QUESTION

BOB KAHN

Uh, yes, in a very material way. Cybelle(?) Weinstein(?) and I actually sat down and wrote the technical part of the proposal that went in. The proposal was packaged as a much larger proposal, and it involved many other people at

BB&N. It was done as part of a group that, uh, Frank Hart(?) led. But Cybelle and I actually put together the technical part of the proposal. I actually have the handwritten version of it that he and I wrote during an all-nighter session at his house.

QUESTION

BOB KAHN

Uh, I thought the proposal that we put together was extremely good. And I thought that it had a really good chance of winning. Uh, but at the time that I wrote it, I don't actually think that, uh, I had understood that I ... I really wanted to be part of the implementation of that activity. I was so thinking seriously about staying with research and in fact going back to MIT to become a professor there again. But, uh, in the final analysis, I thought that it really needed somebody to stay with the system design, and so I did.

QUESTION

BOB KAHN

Yeah, there was a tremendous amount of excitement in the group when the award was won. I was sort of 180 degrees out of phase with it. I ... I was really excited when the proposal was written, because we got our ideas down on paper. And when the contract actually came in, I was kind of in a low, because, uh, I knew at the point the work actually had to be done.

QUESTION

BOB KAHN

You mean after the contract was awarded?

QUESTION

BOB KAHN

No comment on that.

QUESTION

BOB KAHN

The Washington Hilton?

QUESTION

BOB KAHN

Uh, in many ways, the unsung hero of that 1972 demonstration is Al Buzza(?) from MIT. Al never really got the visible recognition, but he and I worked very, very closely in putting the demonstration together. It involved bringing together, oh, some 40 terminal manufacturers, establishing an Arpinet(?) note in the ... in the Hilton. And in fact, the timetable on which we actually had to do the installation was something like six hours to get the phone lines in, to get the machines up before the thing actually opened up. And Al just did a tour de force in managing the whole crew to do that. Uh, it took us about, I would say ... I would say, close to a year to actually plan and carry out that demonstration. And the significance of the demonstration was that it was the first really forceful push to make the Arpinet useful to the end users. Uh, in

1969, we installed the first net, the first net node(?) on the Arpinet at UCLA. And by December of that year we had the first four nodes in place. And Dave Walden and I went out to demonstrate the workings of the network, or not, as the case may be, and shortly thereafter. But by the mid-197Ø, one timeframe, there was still not very much use of the net. It could move packets around, but there weren't machines connected doing useful things. And that project was set up to really force the utility of the Arpinet to occur to the end users. And by that 1972 October demonstration at the Washington Hilton, we had to shoot (Inaudible).

QUESTION

BOB KAHN

No. I think, uh, it's fair to say that we were relatively confident that this could be done, that the community would follow suit. I menton Al particularly because he is a ... really a figure that

needs the recognition for his efforts there. But it was really a community-wide effort. We had at least, uh, 30 or 40 significant major people from the field, from the academic and the research community that contributed to making that a success. People whose names are household words today were the workers in the field. It was just a marvelous team effort all around.

QUESTION

BOB KAHN

I think most of the attendees that had been part of the Arpinet experience at that time obviously knew what this was about. But it was exciting because it was the first time the whole community had shown up in one place at one time. If somebody had dropped a bomb on the Washington Hilton, it would've destroyed almost all of the networking community in the United States at that point. On the other hand, the people who attended that knew very little about

networking Arpinet, I think, in many cases, were astonished about, not only what it could do, but the fact that it worked at all. I really put packet-switching on the map, in my opinion.

QUESTION

BOB KAHN

Well, when I first started working on the whole concept of packet-switching, which was really in the 1966 timeframe, my view was that this was an interesting way to do computer communications. I think by the time the Arpinet project was underway, it was pretty clear to me that this was actually a practical, viable alternative for the marketplace. But whether it would accept it or not, I did not know. I think after the demonstration at the Washington Hilton, it was pretty clear to me that this was gonna fly commercially as well. Of course, Steve Levy(?) and I were very active in putting together the early plans for Telenet, the first

commercial packet net that, uh, made it in the marketplace. And, uh, I think that was one more piece of evidence that this was likely to fly, when, in fact, Telenet made it.

QUESTION

BOB KAHN

Well, the Arpinet was really fairly tightly controlled as a network activity. It was run, policy-wise, out of ARPA(?). It was really run technically out of ARPA from the early days as well, with BB&N playing the key contractor role for the subnet. Uh, but I think the thing you're referring to really came about because of another development, which was really brought upon by the existence of other networking technologies, not just the Arpinet, that led to the creation of what is now known as the Immernet(?). And that's what's had the broad reach around the world. It was not the Arpinet ...

QUESTION

BOB KAHN

I mean, it's gonna be the notion of someone asking the questions?

QUESTION

BOB KAHN

Okay, right.

QUESTION

BOB KAHN

Okay.

QUESTION

BOB KAHN

Well, the Arpinet itself was a network that was fairly tightly managed by ARPA, certainly during its early years and subsequent to that. In the early 1970s, however, I was involved in the, uh, development of several other networking technologies, based upon the same ideas that went into the Arpinet. One of these was a ground based packet radio system that had all of its nodes completely mobile, using

microcomputers and spread spectrum. Uh, the second and ... second network that we were involved with was a network that eventually was implemented on the Intel Sec(?) IV satellite. Called Satnap(?). And those three networks seemed to me, uh, were going to need to be interconnected if we needed to get ... let's try this one again, I'm not happy ...

QUESTION

BOB KAHN

That's ... yeah, let's go back. I actually had a different train of thought I wanted to ...

QUESTION

BOB KAHN

Yep. Uh, the Arpinet was, uh, a relatively small-scale network by comparison with, let's say, the Internet. It was controlled by the Advance Research Projects Agency in a fairly tight way for most of its existence. In order to get on that net, you really needed to either be a defense

contractor, or be approved ... the government paid for all the connections. When you compare that with the Internet, which had a very broad reach ultimately, and reached around the world, I think the real implications of this are being felt more from the Internet than they are from the Arpinet per se. As far as the origins of the Internet, it really goes back to the early 1970s, uh, when, in addition to the Arpinet I was in the process of designing two other networks, one of which was a ground-base packet radio network that embodied many of the notions of the Arpinet package switching concept that applied to small, mobile nodes that, uh, used radio propagation. Uh, the other was a satellite-based net called Satnet, that used Intel Sat IV. It turned out that BB&N was involved in ... in both of those, but in the ground radio network, there were approximately 17 contractors that were involved, and it was managed out of ARPA,

under my auspices. Uh, in the case of the satellite net, that actually started under Larry Roberts' tenure at ARPA, and I'd gotten involved in it. Uh, and it turned out to be a network that involved multiple countries in Europe and the United State on the Intel Sat IV satellite. I knew that for computers on one network to talk to computers on another network, we would have to understand how to link the networks together and how to make it possible for computers on both ends to talk to each other. That meant developing a methodology to do that, because these networks were all different, they had different data rates, they had different packet formats, they had different error structures, error controls, different, uh, methodologies. And that's what led to the develop of the TCPIP protocol suite as a method for that. This is an effort that, uh, Vint Surf(?) and I collaborated on. Vint had been originally involved in the early NCP

efforts, along with, uh, Steve Crocker(?) and ... and a number of other people. And so Vint was very knowledgeable about both the original ~~Pros~~ Host Protocol(?) development for the Arpinet, and the requirements of the machines to house it. I had a pretty good idea of how this ought to work, but I was looking at it from the communications point of view and not necessarily from the needs of the host machines. And Vint and I got together, pooled our talents and the result was the TCPIP Protocol Suite, which in essence formed the framework for the Internet.

QUESTION

BOB KAHN

I don't think there's any single characterization of what it was like to be a pioneer in developing the Arpinet. In fact, there were so many people involved in it that I remember one article where people said that Arpa didn't even know what it was doing, because one article said it was trying

to develop packet switching, and another article said it was trying to do computer resource sharing and couldn't make up its own mind. In fact, the ... the essential element of a really good ARPA program was it was solving multiple problems, achieving multiple objectives, often simultaneously. This was an example where the community that was involved was in fact very large. It involved people across the country in various disciplines, all contributing in their own little way. And I think if you were to write the history of the Arpinet, it would be pretty hard to do a complete job referencing every single person who made a contribution to it.

QUESTION

BOB KAHN

Well, it depends on ... the essence of the Arpinet development really depends on what you think of as the Arpinet. To some people, they think about it as just the communication pieces

of the network. Other pieces (sic) think about, you know, the use of Jabet(?). Other people think about the strategy and the structure and all of that sort of stuff. And I think those were all different communities, all interacting in interesting ways, but there was no quintessential Arpinet team when you look at it broadly.

QUESTION

BOB KAHN

There was no question in my mind that when they installed that in, that it would do what we expected to do, which was, one, end function, and send packets out the lines and bring packets in. You remember, one node network is about as useful as one telephone system. So the best you could do was loop some lines and make sure that what went out a line came back in the same line. Or that the software didn't crash, or that the internal activities were functioning properly. There was not a major effort to try and check

out that installation when the first node was installed, because there wasn't a whole lot to do, other than to be sure the software ran. It was after the fourth node was installed ... we actually had a four-node subnet on the West Coast, that Dave Walden and I went out to the West Coast to actually do the field testing and debugging of the system.

QUESTION

BOB KAHN

Nope.

QUESTION

BOB KAHN

What's the if?

QUESTION

BOB KAHN

Well, of course, the Arpinet was decommissioned by ARPA in 1990. And the Internet is gradually becoming more and more embraced by commercial industry. I think the

basic notion of the Arpinet was fundamentally sound, in terms of the packet-switching approach and the ability to break things into small trunks and send them, I think the fundamental notion of the Internet was sound, in terms of linking independent systems without requiring central and overriding the administration of everything.

In 1985, I left ARPA specifically to work on fostering the development of the National Information Infrastructure in the United States. At the time, this was not a well-known term, and I think many people probably viewed it as equivalent to harvesting the clouds. But over the course of the last eight years, there's been a growing awareness of the importance of this whole activity, and the importance not only to this country, but perhaps to the ... to the whole world. The Internet plays a very important role in that it currently is as close to an infrastructural base for the global infrastructure as I think exists

in the world today. And the credit for that goes to a lot of the people that made that possible around the world. But I think that the future is really going to see the whole notion of packet switching embodied in these national infrastructures, and the global infrastructure. And what we'll see is better, more rapid access to the information that we need, more economic ways to communicate with people, better ways to integrate data and other kinds of information in ways that can benefit all of society.

QUESTION

BOB KAHN

Well, I'm an optimist when it comes to technology and society. I think that almost any technology can be put to positive use. On the converse, almost any technology can be misused and even technologies that have inherently positive uses can be used by small fractions of the population. So I think we have here an

opportunity, and I think it's one that we're going to capitalize on and it's one that, uh, you know, I'm basically very positive about.

QUESTION

BOB KAHN

Well, you know, I think the most memorable aspects of that whole experience for me were the facility when it actually got into existence, partly in the activity that Dave Walden and I had when we went out to try to bug the first instants of it. The feel of satisfaction ... I saw ... actually the first public demonstration of it at the Washington Hilton. And I think if I had to pick any singular event during that period, it would probably be the efforts that Len Klein, Arquel(?), Frank and I had in actually writing the first paper on lessons learned in the Arpinet experience. That was, uh ... it was a wonderful interchange between three people who were then and are still very good friends. But actually brought

completely different points of view to the same common experience we had all been through. Len had been approaching this whole activity from the point of view of the theorist. He had an analytical notion. I mean, his view was that, you know, we can build networks up to about 40 to 60 nodes. And Howard and I both shared that view. But Len's argument was that the denominator of some equation went to zero. And Howard would say, from a simulation point of view, the simulations don't work beyond that point. And I would say, there's something in the implementation, like the routing cables, that we got too big on the lines, or something beyond that ... that size. So the interplay of our three different approaches was something that was really fun to work through and to actually sit down and write.

QUESTION

BOB KAHN

Well, Howard ... Howard and I actually got together once, I recall, up in Boston. And it's possible that Len and I actually chatted on one or more occasions at his office at UCLA, although I have no specific recollection of that. Uh, I do have a very specific recollection that we did not all get together at any one point at any one time, so there surely was no pizza that we all shared simultaneously. But maybe Len and Howard will contradict that.

QUESTION

BOB KAHN

Remember, it was 20 years ago.

QUESTION

BOB KAHN

(END OF TAPE 3)